arXiv:1612.02336v1 [cs.NE] 7 Dec 2016

Neural Turing Machines: Convergence of Copy Tasks

Janez Aleš
Heidelberg Colaboratory for Image Processing
IWR, Heidelberg University
69120 Heidelberg, Germany
janez.ales@iwr.uni-heidelberg.de

December 8, 2016

Abstract

The architecture of neural Turing machines is differentiable end to end and is trainable with gradient descent methods. Due to their large unfolded depth Neural Turing Machines are hard to train and because of their linear access of complete memory they do not scale. Other architectures have been studied to overcome these difficulties. In this report we focus on improving the quality of prediction of the original linear memory architecture on copy and repeat copy tasks. Copy task predictions on sequences of length six times larger than those the neural Turing machine was trained on prove to be highly accurate and so do predictions of repeat copy tasks for sequences with twice the repetition number and twice the sequence length neural Turing machine was trained on.

1 Introduction

Neural network capabilities were extended by coupling them to linear external memory accessed by probability distribution weight vectors, called Neural Turing Machines (NTM) [1]. This architecture is differentiable end to end and can be trained with gradient descent. It is known that their large unfolded depth makes them hard to train. They also do not scale well due to the linear access of their complete memory. Alternate memory architectures have since been proposed to improve training behaviour. Structured memory components for neural Turing machines were tested for speed of convergence and quality of predictions in [4]. In contrast to the original architecture neural GPUs are higly parallelizable and scale well [3].

In this report we focus on the quality of predictions of the original linear memory architecture of NTMs (see [1]) for copy and repeat copy tasks on *longer* sequences than those the neural turing machine was trained on. Prediction on these sequences prove to be accurate.

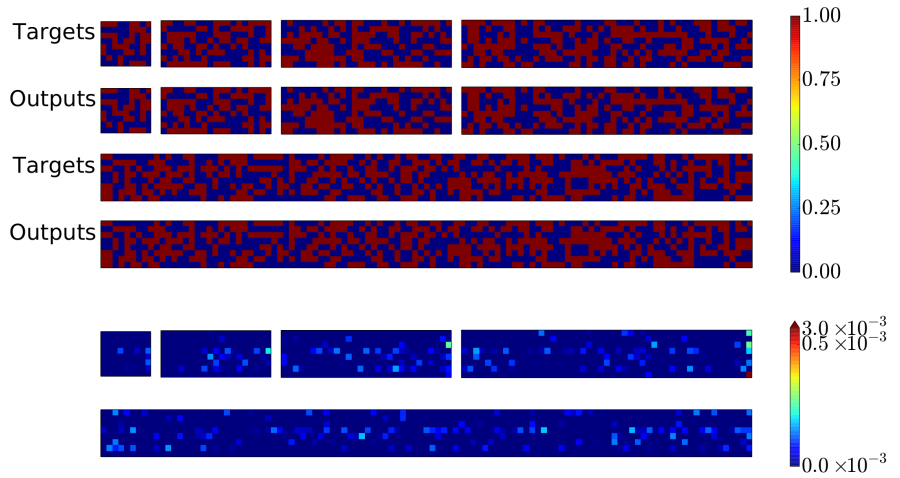

Figure 1: Targets and outputs for sequences of length 10, 20, 30, 50, and 120, followed by differences of outputs and targets with a logarithmic colormap.

2 Copy Task

An NTM with a single layer feedforward controler of size 100 and external memory size 128×20 is trained for a copy task on sequences of random binary vectors of length 8 with sequence lengths chosen randomly between 1 and 20. Prediction quality for this task was already examined in [1]. Tests on sequences of legth 10 and 20 reproduced inputs with high confidence and with virtually no mistakes. However, predictions on sequences of length 30 and 50 show a few errors on a relatively small number of bits. Predictions on sequences of length 120 produce local and global errors leading to high loss. A higher accuracy on the copy task than in the original paper was achieved in [2] using structured memory NTMs.

We present a solution using the original NTM architecture. Learning stage was carried out on sequences of random length up to 20 with the sum of bit values *not* equal to half of the number of bits. Testing was done on the complement set, that is, on the sequences with the sum of bit values *equal* to half of the number of bits assuring independent test results. Predictions reproduced inputs with high confidence for sequences of length 10, 20, 30, 50, and 120. Prediction statistics were calculated on the basis of 10000 random instances for each sequence length. We discuss the results of a representative set. Predictions with no bit errors are shown in Figure 1. Sequence lengths, number of sequences with at least one bit error, maximum bit error, mean of bit errors, and standard deviation of bit errors are shown in the Table 1. The maximum bit error was 1 and there were no sequences with a global error. Learning curve for the copy task is shown in Figure 3.

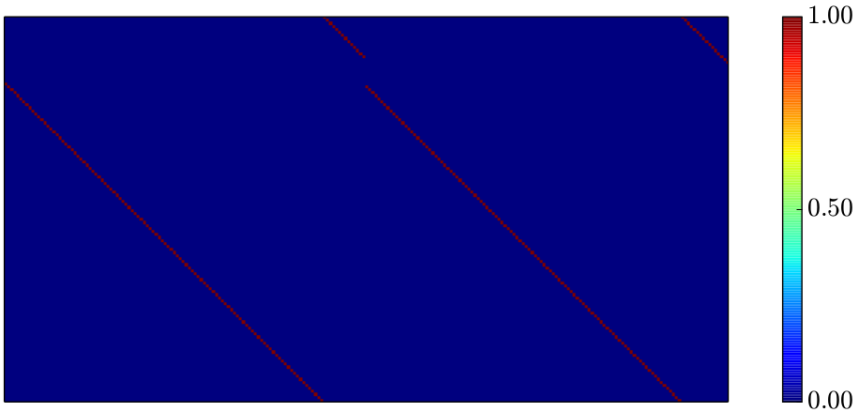

Figure 2: Write and read weightings for a sequence of length 120.

Table 1: Prediction statistics for a copy task on a representative set of 10000 sequences.

Sequence length	10	20	30	50	120
Number of seq. with bit errors	0	0	0	13	36
Maximum bit error	0	0	0	1	1
Mean of bit errors	0.0000	0.0000	0.0000	0.0013	0.0036
Standard deviation of bit errors	0.0000	0.0000	0.0000	0.0360	0.0599

3 Repeat Copy Task

As in the original paper we train the NTM on 6 bit vector sequences of length up to 10 and repetition number up to 10. We test the quality of the prediction on sequences of length 10 with repetition number 20, on sequences of length 20 with repetition number 10, and on sequences of length 20 with repetition number 20. Prediction statistics was based on 10000 random sequences. We discuss a representative result. [1]

Prediction for sequences of length 20 with repeat number 10 had 78 sequences having at least one faulty bit. Maximum bit error was 7 with bit error mean of 0.013 and bit error standard deviation of 0.183.

Prediction for sequences of length 20 with repeat number 20 had 104 sequences having at least one faulty bit. Maximum bit error was 10 with bit error mean 0.018 and bit error standard deviation of 0.254.

Prediction on sequences of length 10 with repeat number 20 were less accurate

[1] Predictor weights can be provided upon request.

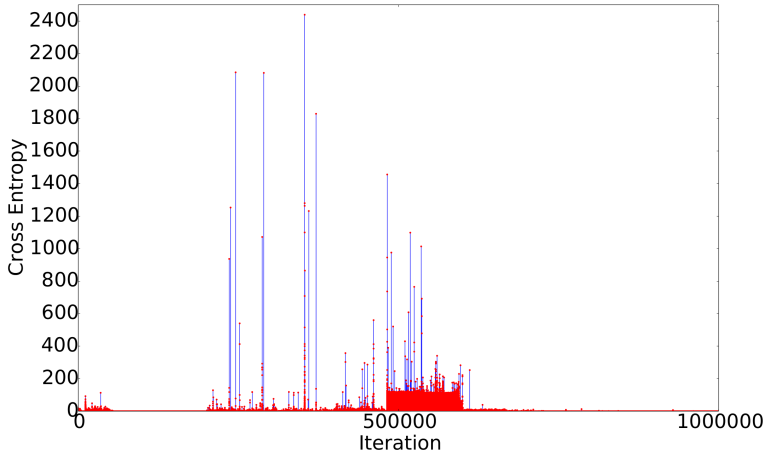

Figure 3: Copy learning curve.

with bit error mean 2.449 and bit error standard deviation of 8.391. In total 4418 sequences had at least one bit error and 19 sequences had more than 100 faulty bits.

Alternate NTM training sessions of repeat copy task will be addressed in future work.

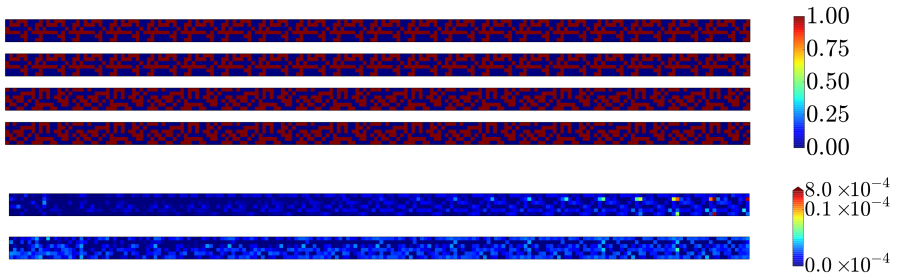

Figure 4: Targets and outputs of the repeat copy task for sequences of length 10 with repetition number 20 and length 20 with repetition number 10, followed by the corresponding differences.

Figure 5: Write and read weightings for the repeat copy task for sequences of length 10 with repetition number 20 and length 20 with repetition number 10.

Figure 6: Write and read weightings for the repeat copy task for a sequence of length 10 and repetition number 20 with a global error.

References

[1] Alex Graves, Greg Wayne, and Ivo Danihelka. Neural Turing Machines. *arXiv*, 1410(5401v2):1–26, 2014.

[2] Rasmus Boll Greve, Emil Juul Jacobson, and Sebastian Risi. Evolving Neural Turing Machines. *Neural Information Processing Systems: Reasoning, Attention, Memory Workshop*, 2015.

[3] Lukasz Kaiser and Ilya Sutskevar. Neural GPUs Learn Algorithms. *International Conference on Learning Representations*, 2016.

[4] Wei Zhang, Yang Yu, and Bowen Zhou. Structured Memory for Neural Turing Machines. *arXiv*, 1510(03931v3):1–4, 2015.

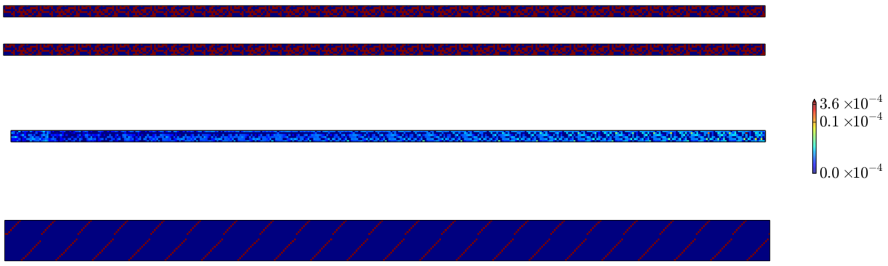

Figure 7: Targets and outputs of the repeat copy task for sequences of length 20 and repetition number 20. The corresponding differences are shown in logarithmic scale, followed by write and read weightings.